

 **PORTAL**  
USPTO

Subscribe (Full Service) Register (Limited Service, Free) Login  
 Search:  The ACM Digital Library  The Guide  
 +optimize +query +view +"join predicate" +"base table" +uni **SEARCH**

 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [optimize](#) [query](#) [view](#) [join](#) [predicate](#) [base table](#) [union](#) [false](#) [empty](#)

Found 27 of 185,030

Sort results by

relevance 

 [Save results to a Binder](#)

Try an [Advanced Search](#)

Display results

expanded form 

 [Search Tips](#)

Try this search in [The ACM Guide](#)

Open results in a new window

Results 1 - 20 of 27

Result page: **1** [2](#) [next](#)

Relevance scale 

### 1 [Extensible/rule based query rewrite optimization in Starburst](#)



Hamid Pirahesh, Joseph M. Hellerstein, Waqar Hasan  
 June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data SIGMOD '92**, Volume 21 Issue 2

Publisher: ACM Press

Full text available:  [pdf\(1.30 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes the Query Rewrite facility of the Starburst extensible database system, a novel phase of query optimization. We present a suite of rewrite rules used in Starburst to transform queries into equivalent queries for faster execution, and also describe the production rule engine which is used by Starburst to choose and execute these rules. Examples are provided demonstrating that these Query Rewrite transformations lead to query execution time improvements of orders of magni ...

### 2 [Research session: views and cache management: View matching for outer-join views](#)

Per-Åke Larson, Jingren Zhou

August 2005 **Proceedings of the 31st international conference on Very large data bases VLDB '05**

Publisher: VLDB Endowment

Full text available:  [pdf\(279.76 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Prior work on computing queries from materialized views has focused on views defined by expressions consisting of selection, projection, and inner joins, with an optional aggregation on top (SPJG views). This paper provides the first view matching algorithm for views that may also contain outer joins (SPOJG views). The algorithm relies on a normal form for SPOJ expressions and does not use bottom-up syntactic matching of expressions. It handles any combination of inner and outer joins, deals cor ...

### 3 [Simplification of outer joins](#)

Gautam Bhargava, Piyush Goel, Balakrishna R. Iyer

November 1995 **Proceedings of the 1995 conference of the Centre for Advanced Studies on Collaborative research**

Publisher: IBM Press

Full text available:  [pdf\(309.09 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The removal of redundant outer joins is essential for the reassociation of outer joins with

other binary operations. In this paper, we present a set of comprehensive algorithms that employ the properties of strong predicates along with the properties of SQL's projection, intersection, union and except operations in order to remove redundant outer joins from a complex query. For the purpose of query simplification, we generate additional projections by determining the keys. Our algorithm for gene ...

#### **4 Fundamental techniques for order optimization**

 David Simmen, Eugene Shekita, Timothy Malkemus  
June 1996 **ACM SIGMOD Record , Proceedings of the 1996 ACM SIGMOD international conference on Management of data SIGMOD '96**, Volume 25 Issue 2.

**Publisher:** ACM Press

Full text available:  pdf(1.07 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Decision support applications are growing in popularity as more business data is kept online. Such applications typically include complex SQL queries that can test a query optimizer's ability to produce an efficient access plan. Many access plan strategies exploit the physical ordering of data provided by indexes or sorting. Sorting is an expensive operation, however. Therefore, it is imperative that sorting is optimized in some way or avoided all together. Toward that goal, this paper describe ...

#### **5 A formal perspective on the view selection problem**

Rada Chirkova, Alon Y. Halevy, Dan Suciu

November 2002 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 11 Issue 3

**Publisher:** Springer-Verlag New York, Inc.

Full text available:  pdf(329.63 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

The view selection problem is to choose a set of views to materialize over a database schema, such that the cost of evaluating a set of workload queries is minimized and such that the views fit into a prespecified storage constraint. The two main applications of the view selection problem are materializing views in a database to speed up query processing, and selecting views to materialize in a data warehouse to answer decision support queries. In addition, view selection is a core problem for i ...

**Keywords:** Materialized views, View selection

#### **6 Efficiently publishing relational data as XML documents**

Jayavel Shanmugasundaram, Eugene Shekita, Rimon Barr, Michael Carey, Bruce Lindsay, Hamid Pirahesh, Berthold Reinwald

September 2001 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 10 Issue 2-3

**Publisher:** Springer-Verlag New York, Inc.

Full text available:  pdf(216.67 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

XML is rapidly emerging as a standard for exchanging business data on the World Wide Web. For the foreseeable future, however, most business data will continue to be stored in relational database systems. Consequently, if XML is to fulfill its potential, some mechanism is needed to publish relational data as XML documents. Towards that goal, one of the major challenges is finding a way to efficiently structure and tag data from one or more tables as a hierarchical XML document. Different alterna ...

**Keywords:** Publishing, Relational databases, XML

#### **7**

#### **Optimizing queries using materialized views: a practical, scalable solution**



Jonathan Goldstein, Per-Åke Larson

May 2001 **ACM SIGMOD Record , Proceedings of the 2001 ACM SIGMOD international conference on Management of data SIGMOD '01**, Volume 30 Issue 2

Publisher: ACM Press

Full text available: pdf(202.08 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Materialized views can provide massive improvements in query processing time, especially for aggregation queries over large tables. To realize this potential, the query optimizer must know how and when to exploit materialized views. This paper presents a fast and scalable algorithm for determining whether part or all of a query can be computed from materialized views and describes how it can be incorporated in transformation-based optimizers. The current version handles views composed of sele ...

**Keywords:** materialized views, query optimization, view matching

## 8 The state of the art in distributed query processing



Donald Kossmann

December 2000 **ACM Computing Surveys (CSUR)**, Volume 32 Issue 4

Publisher: ACM Press

Full text available: pdf(455.39 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Distributed data processing is becoming a reality. Businesses want to do it for many reasons, and they often must do it in order to stay competitive. While much of the infrastructure for distributed data processing is already there (e.g., modern network technology), a number of issues make distributed data processing still a complex undertaking: (1) distributed systems can become very large, involving thousands of heterogeneous sites including PCs and mainframe server machines; (2) the stat ...

**Keywords:** caching, client-server databases, database application systems, dissemination-based information systems, economic models for query processing, middleware, multitier architectures, query execution, query optimization, replication, wrappers

## 9 Multidatabase systems: Exploiting uniqueness in query optimization



G. N. Pauley, Per-Åke Larson

October 1993 **Proceedings of the 1993 conference of the Centre for Advanced Studies on Collaborative research: distributed computing - Volume 2**

Publisher: IBM Press

Full text available: pdf(1.27 MB)

Additional Information: [full citation](#), [abstract](#), [references](#)

Consider an SQL query that specifies duplicate elimination via a DISTINCT clause. Because duplicate elimination often requires an expensive sort of the query result, it is often worthwhile to identify situations where the DISTINCT clause is unnecessary, to avoid the sort altogether. We prove a necessary and sufficient condition for deciding if a query requires duplicate elimination. The condition exploits knowledge about keys, table constraints, and query predicates. Because the condition cannot ...

## 10 Research papers: adaptive, automatic, autonomic systems: Automatic physical database tuning: a relaxation-based approach



Nicolas Bruno, Surajit Chaudhuri

June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: [pdf\(476.05 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

In recent years there has been considerable research on automated selection of physical design in database systems. In current solutions, candidate access paths are heuristically chosen based on the structure of each input query, and a subsequent bottom-up search is performed to identify the best overall configuration. To handle large workloads and multiple kinds of physical structures, recent techniques have become increasingly complex: they exhibit many special cases, shortcuts, and heuristics ...

**11 Research papers: optimization: RankSQL: query algebra and optimization for relational top-k queries**



Chengkai Li, Kevin Chen-Chuan Chang, Ihab F. Ilyas, Sumin Song  
June 2005 **Proceedings of the 2005 ACM SIGMOD international conference on Management of data**

Publisher: ACM Press

Full text available: [pdf\(741.54 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

This paper introduces RankSQL, a system that provides a systematic and principled framework to support efficient evaluations of ranking (*top-k*) queries in relational database systems (RDBMS), by extending relational algebra and query optimization. Previously, *top-k* query processing is studied in the middleware scenario or in RDBMS in a "piecemeal" fashion, *i.e.*, focusing on specific operator or sitting outside the core of query engines. In contrast, we aim to support ranking ...

**12 Extensible query processing in starburst**



L. M. Haas, J. C. Freytag, G. M. Lohman, H. Pirahesh  
June 1989 **ACM SIGMOD Record , Proceedings of the 1989 ACM SIGMOD international conference on Management of data SIGMOD '89**, Volume 18 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.63 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Today's DBMSs are unable to support the increasing demands of the various applications that would like to use a DBMS. Each kind of application poses new requirements for the DBMS. The Starburst project at IBM's Almaden Research Center aims to extend relational DBMS technology to bridge this gap between applications and the DBMS. While providing a full function relational system to enable sharing across applications, Starburst will also allow (sophisticated) programmers to add many kinds of ...

**13 Answering complex SQL queries using automatic summary tables**



Markos Zaharioudakis, Roberta Cochrane, George Lapis, Hamid Pirahesh, Monica Urata  
May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data SIGMOD '00**, Volume 29 Issue 2

Publisher: ACM Press

Full text available: [pdf\(185.85 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We investigate the problem of using materialized views to answer SQL queries. We focus on modern decision-support queries, which involve joins, arithmetic operations and other (possibly user-defined) functions, aggregation (often along multiple dimensions), and nested subqueries. Given the complexity of such queries, the vast amounts of data upon which they operate, and the requirement for interactive response times, the use of materialized views (MVs) of similar complexity is often mandatory ...

**14 On the problem of generating common predecessors**



W. Lehner, W. Hüpper, L. Schlesinger, A. Bauer  
November 2002 **Proceedings of the 5th ACM international workshop on Data**

**Warehousing and OLAP****Publisher:** ACM PressFull text available:  pdf(231.16 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Using common subexpressions to speed up a set of queries is a well known and long studied problem. However, due to the isolation requirement, operating a database in the classic transactional way does not offer many applications to exploit the benefits of simultaneously computing a set of queries. In the opposite, many applications can be identified in the context of data warehousing, e. g. optimizing the incremental maintenance process of multiple dependent materialized views or the generation ...

**15 Extensions to Starburst: objects, types, functions, and rules**
 Guy M. Lohman, Bruce Lindsay, Hamid Pirahesh, K. Bernhard Schiefer  
October 1991 **Communications of the ACM**, Volume 34 Issue 10**Publisher:** ACM PressFull text available:  pdf(5.21 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

**Keywords:** Extended relational database management systems, Starburst, extensible database management systems

**16 Active rules in deductive databases**
 John V. Harrison  
December 1993 **Proceedings of the second international conference on Information and knowledge management****Publisher:** ACM PressFull text available:  pdf(1.13 MB) Additional Information: [full citation](#), [references](#), [index terms](#)**17 Cost-based optimization of decision support queries using transient-views**
 Subbu N. Subramanian, Shivakumar Venkataraman  
June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98**, Volume 27 Issue 2**Publisher:** ACM PressFull text available:  pdf(1.58 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Next generation decision support applications, besides being capable of processing huge amounts of data, require the ability to integrate and reason over data from multiple, heterogeneous data sources. Often, these data sources differ in a variety of aspects such as their data models, the query languages they support, and their network protocols. Also, typically they are spread over a wide geographical area. The cost of processing decision support queries in such a setting is quite high. Ho ...

**18 Parallelism and its price: a case study of nonstop SQL/MP**
 Susanne Englert, Ray Glasstone, Waqar Hasan  
December 1995 **ACM SIGMOD Record**, Volume 24 Issue 4**Publisher:** ACM PressFull text available:  pdf(1.09 MB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

We describe the use of parallel execution techniques and measure the price of parallel execution in NonStop SQL/MP, a commercial parallel database system from Tandem Computers. Non-Stop SQL uses intra-operator parallelism to parallelize joins, groupings and scans. Parallel execution consists of starting up several processes and communicating

data between them. Our measurements show (a) Startup costs are negligible when processes are reused rather than created afresh (b) Communication costs ...

## **19 Optimization techniques for queries with expensive methods**



Joseph M. Hellerstein

June 1998 **ACM Transactions on Database Systems (TODS)**, Volume 23 Issue 2

Publisher: ACM Press

Full text available: pdf(582.16 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Object-relational database management systems allow knowledgeable users to define new data types as well as new methods (operators) for the types. This flexibility produces an attendant complexity, which must be handled in new ways for an object-relational database management system to be efficient. In this article we study techniques for optimizing queries that contain time-consuming methods. The focus of traditional query optimizers has been on the choice of join methods and orders; selec ...

**Keywords:** expensive methods, extensibility, object-relational databases, predicate migration, predicate placement, query optimization

## **20 Paper session DB-9 (databases): query processing 1: Selectivity-based partitioning:**



**a divide-and-union paradigm for effective query optimization**

Neoklis Polyzotis

October 2005 **Proceedings of the 14th ACM international conference on Information and knowledge management CIKM '05**

Publisher: ACM Press

Full text available: pdf(253.90 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Modern query optimizers select an efficient join ordering for a physical execution plan based essentially on the average join selectivity factors among the referenced tables. In this paper, we argue that this "monolithic" approach can miss important opportunities for the effective optimization of relational queries. We propose *selectivity-based partitioning*, a novel optimization paradigm that takes into account the join correlations among *relation fragments* in order to essen ...

Results 1 - 20 of 27

Result page: **1** [2](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2006 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#)

[QuickTime](#)

[Windows Media Player](#)

[Real Player](#)

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

[Search Results](#)[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Results for "((join predicate' &lt;and&gt; query &lt;and&gt; view &lt;and&gt; union)&lt;in&gt;metadata)"

Your search matched 0 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance in Descending order**.» **Search Options**[View Session History](#)[Modify Search](#)[New Search](#)

((join predicate' &lt;and&gt; query &lt;and&gt; view &lt;and&gt; union)&lt;in&gt;metadata)

 Check to search only within this results setDisplay Format:  Citation  Citation & Abstract» **Key**

IEEE JNL IEEE Journal or Magazine

IEE JNL IEE Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IEE CNF IEE Conference Proceeding

IEEE STD IEEE Standard

**No results were found.**

Please edit your search criteria and try again. Refer to the Help pages if you need assistance.

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

Indexed by  
 Inspec®

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

[Search Results](#)[BROWSE](#)[SEARCH](#)[IEEE XPLORE GUIDE](#)

Results for "((join &lt;and&gt; predicate &lt;and&gt; query &lt;and&gt; view &lt;and&gt; union)&lt;in&gt;metadata..."



Your search matched 1 of 1397873 documents.

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance** in **Descending** order.» **Search Options**[View Session History](#)[Modify Search](#)[New Search](#) Check to search only within this results setDisplay Format:  Citation  Citation & Abstract [Select All](#) [Deselect All](#)

1. Efficient processing of outer joins and aggregate junctions  
Bhargava, G.; Goel, P.; Iyer, B.;  
[Data Engineering, 1996. Proceedings of the Twelfth International Conference on](#)  
26 Feb.-1 March 1996 Page(s):441 - 449  
Digital Object Identifier 10.1109/ICDE.1996.492193  
[AbstractPlus](#) | Full Text: [PDF\(676 KB\)](#) [IEEE CNF](#)  
[Rights and Permissions](#)

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

Indexed by

[Return to the USPTO NPL Page](#) | [Help](#)[Basic](#)[Advanced](#)[Topics](#)[Publications](#)[My Research](#)

0 marked items

Interface language:

English

Databases selected: Multiple databases...**No documents found for: (join w/2 predicate and view\* and quer\* and table) AND PDN(<9/24/2003)**Refine your search below using the following tips:

- Check your spelling.
- Reduce the number of terms included in your search.
- Broaden your search by selecting other databases, removing limits, or searching "Citations and document text" (if available).
- Use "AND" to connect two words that don't need to be searched as a phrase.
- Connect similar terms with the "OR" operator (e.g. military OR pentagon). See Search Tips for more hints.

## Basic Search

[Tools: Search Tips](#) [Browse Topics](#) [1 Recent Searches](#)  Database:   [Select multiple databases](#)Date range:   [About](#)Limit results to:  [Full text documents only](#)  [Scholarly journals, including peer-reviewed](#) [More Search Options](#)Copyright © 2006 ProQuest Information and Learning Company. All rights reserved. [Terms and Conditions](#)[Text-only interface](#)

[Return to the USPTO NPL Page](#) | [Help](#)[Basic](#)[Advanced](#)[Topics](#)[Publications](#)[My Research](#)

0 marked items

Interface language:

English

Databases selected: Multiple databases...**No documents found for: (join and predicate and view\* and quer\* and table) AND PDN(<9/24/2003)****Refine your search** below using the following tips:

- Check your spelling.
- Reduce the number of terms included in your search.
- Broaden your search by selecting other databases, removing limits, or searching "Citations and document text" (if available).
- Use "AND" to connect two words that don't need to be searched as a phrase.
- Connect similar terms with the "OR" operator (e.g. military OR pentagon). See Search Tips for more hints.

## Basic Search

[Tools: Search Tips](#) [Browse Topics](#) [2 Recent Searches](#)

Database:

[Select multiple databases](#)

Date range:



09/24/2003

[About](#)

Limit results to:

 Full text documents only  Scholarly journals, including peer-reviewed [About](#)[More Search Options](#)Copyright © 2006 ProQuest Information and Learning Company. All rights reserved. [Terms and Conditions](#)[Text-only interface](#)



August 30, 2006

USPTO

**Search**

- [Full Text](#)
- [Concept](#)
- [Document ID](#)
- [Recent Disclosures](#)

**Other**

- [Prior Art Home](#)
- [Support](#)
- [Logout](#)

**No records matched your search.**

Perhaps you should try a less restrictive query.

**Search query:** "join predicate" and view and query and join and "base table"[New search](#) | [Modify this search](#)

Copyright © 2006 IP.com, Inc. All rights reserved. |



August 30, 2006

USPTO

Secur

**Search**

Full Text  
Concept  
Document ID  
Recent Disclosures

**Other**

Prior Art Home  
Support  
Logout

Displaying records #1 through 1 out of 1

Result # 1 Relevance: 00000

**Heuristic Method for Joining Relational Data Base Tables**

1988-02-01 IPCOM000056768D

English

This invention relates to a method for joining tables responsive to relational queries. (It  
The method is of the "greedy heuristic" type. This type combines features of both broad  
depth-first search strategies. The method iteratively ...

---

**Displaying page 1 of 1** | << FIRST | < BACK | NEXT > | LAST >>

---

**Search query:** "join predicate" and view and query and join

[New search](#) | [Modify this search](#) | [Search within current results](#)

Copyright © 2006 IP.com, Inc. All rights reserved. |